

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.

(12) UK Patent Application (19) GB (11) 2 270 791 (13) A

(43) Date of Publication 23.03.1994

(21) Application No 9316067.9

(22) Date of Filing 03.08.1993

(30) Priority Data

(31) 07949174 (32) 21.09.1992 (33) US

(71) Applicant(s)

The Grass Valley Group Inc

(Incorporated in USA - California)

13024 Bitney Springs Road, Nevada City,
California 95959, United States of America

(72) Inventor(s)

Kevin D Windrem
David Casper
Anna Greco
Scott Murray
Paul O'Connor
Elaine B Zamani

(51) INT CL⁵

G06F 3/06, G11B 20/10

(52) UK CL (Edition M)

G5R RB788 RB789 RGB
G4A AKB1

(56) Documents Cited

EP 0508441 A2 EP 0507552 A2 EP 0503768 A1
EP 0493984 A2 WO 92/01988 A1 WO 90/00280 A1

(58) Field of Search

UK CL (Edition L) G4A AKB1, G5R RGB
INT CL⁵ G06F, G11B
Online databases: WPI

(74) Agent and/or Address for Service

Langner Parry
High Holborn House, 52-54 High Holborn, LONDON,
WC1V 6RR, United Kingdom

(54) Video disk storage array

(57) A digital video recorder 10 employing standard hard disk arrays 12 employs a caching system 14, 16, 18 to enable continuous video data to be supplied to and provided from the standard disk drives which may not be adapted for continuous data stream operation. The system is implemented as an array of disk drives and includes cache management functions for governing disk read and write operations. A predictor enables the cache to predict disk reads and disk writes and to maintain the cache based on the predictions. A redundant data disk 17 and controller 99 enable playback in the event of failure of a drive in the array and also provides improved playback data retrieval time.

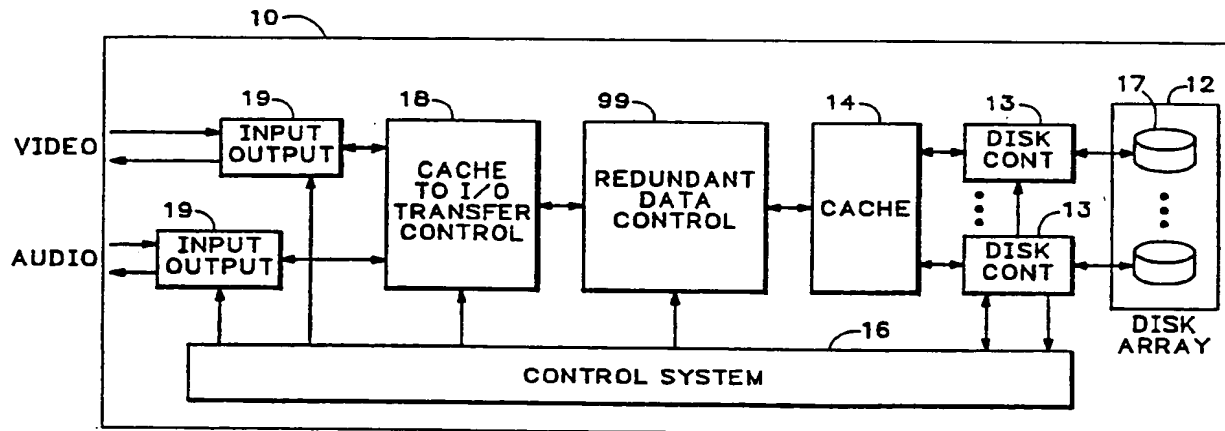


Fig.1

GB 2 270 791 A




Fig.1




Fig. 2


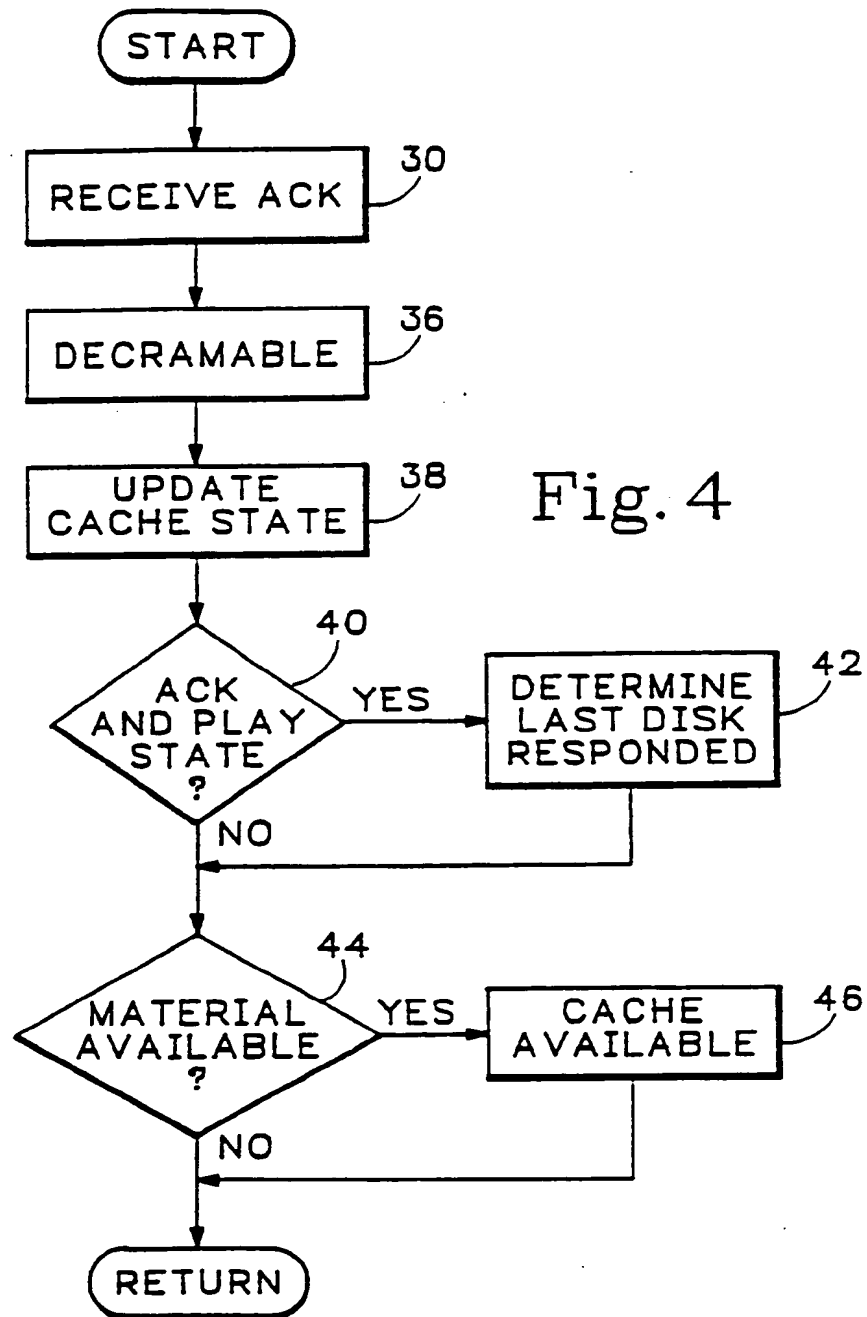


Fig. 3



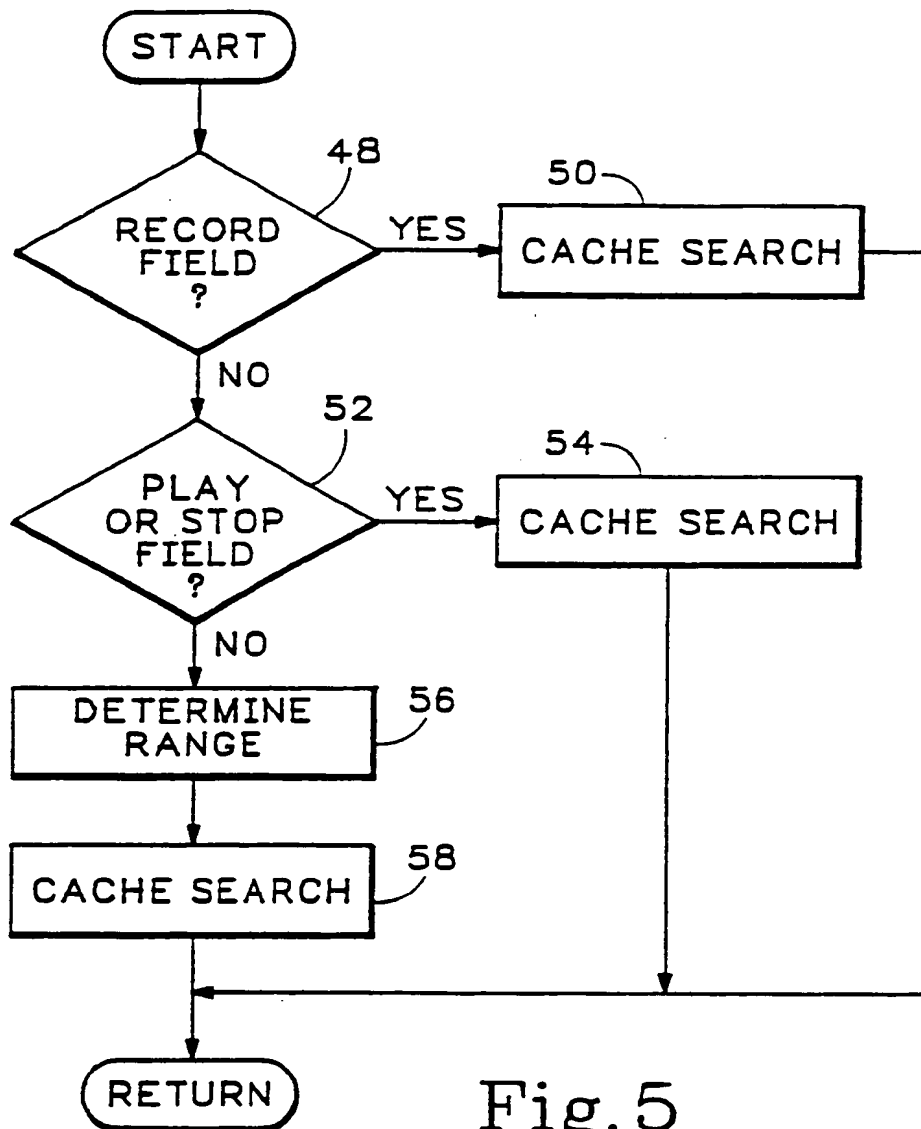


Fig. 5

DISK-BASED DIGITAL VIDEO RECORDER

This invention relates generally to digital video recorders and, more particularly, to cache management and disk utilization for digital video recorders having multiple inputs and outputs.

Background of the invention

An array of disk drives may be assembled to provide sufficient bandwidth to record or play digitized video signals, allowing random access to video data (tape recorders allow only sequential access). However, the data transfer rate of standard disk drives is not constant in that the disk read/write heads must often move from one track to the next and during these periods of movement, no transfer of data to or from the disk may take place. Such limitations are easily overlooked in computer applications since the computer may wait during the periods when no access is available. However, such wait periods are not appropriate in video applications given the continuous stream of data associated with a video signal. While the disks cannot support a continuous transfer data rate required for video, disk arrays are typically able to transfer data in bursts (groups of data) at rates slightly higher than the required continuous rate.

Redundant data storage may be provided to insure accurate data reproduction in case of disk drive failure. However, such redundancy can greatly increase the cost of the system when expensive disk drives are employed. Once a determination of drive failure is made, the

redundant data may be retrieved, but a redundant system typically will wait a relatively long time period before determining that a drive has failed. Such waits may not be appropriate during playback of continuous data.

Another problem attendant with an array of disks is that data retrieval speed is dictated by the slowest disk in the array.

10 Summary of the invention

Cache management/disk utilization for a digital video recorder according to the present invention employs a cache in conjunction with an array of inexpensive disk drives for recording and playing video signals. The cache is such that data coming in or going out (e.g., digital video) is buffered in the cache through a cache management algorithm, thereby compensating for the periods when the disks may be unable to transfer data. Data is maintained in video and audio blocks, enabling use and modification of either or both of video data and audio data. In playback mode, e.g., the cache management includes retrieving playback data from the disks in advance of their use, and storing this information in the cache, and making data available for output.

Brief description of the drawings

For a better understanding of the invention, and to show how the same may be carried into effect, reference will now be made, by way of example, to the accompanying drawings in which:

FIG. 1 is a block diagram of a disk-based digital video recorder according to the present invention;

FIG. 2 is an illustration of the organization of data on the disk array and in the cache according to the present invention;

5 FIG. 3 is a diagram of the various tag states of cache blocks according to the present invention;

FIG. 4 is a flow chart of the Validator process employed by the cache management system; and

10 FIG. 5 is a flow chart of the Query process employed by the cache management system.

Detailed description

FIG. 1 is a block diagram of a digital video recorder employing cache management/disk utilization according to the present invention. The recorder 10 comprises a number of input/output data channels 19 which are supplied through cache to I/O transfer control 18 to a cache 14. Two I/O data channels, one for video data and one for audio data are shown in FIG. 1. Disk array 12, comprising a plurality of disk drives, receives record data and supplies playback data via the cache 14. Operation of the disk array and cache to I/O transfer control is directed by control system 16, which is operative to provide typical video recorder functions, e.g., play, stop, record. The disk drives in the disk array operate asynchronously, and therefore, one drive may complete an operation before another drive completes its corresponding operation.

30 Data is spread across the plurality of drives, providing greater bandwidth than would be available with a single drive. Redundancy is provided through a redundant data controller 99 to handle possible failure of one drive in the array, by exclusive ORing corresponding bits sent to each drive, e.g.,

35

all bits 0 are exclusive OR d generating redundant bit 0. During write operations, this produces the value stored in a redundant drive 17. The data bit value from a non-responding drive may then be reconstructed during a read by the process described above substituting the redundant drive's data for the missing data. The non-responding drive is the last drive to complete its operation.

5
10
15
20
Coupled to disk controllers 13, redundant data control 99, transfer control 18 and I/O data channels 19 is the control system 16 which functions to direct management and operation of the recorder, governing where data is stored and the like. The control system implements stop, play, record, etc. commands, managing the cache and generating disk read and write requests as needed. The control system also provides a high level user interface and an interface to other controlling devices. Primary input to the control system is stimulus from an operator of the system.

25
30
35
All disk controllers receive the same instructions from the control system. Cache 14 consists of multiple dual-ported memories, one memory for each disk drive in disk array 12, each memory having a DMA controller. From the point of view of each disk controller, its portion of memory within the cache is independent of the other controller cache memory portions. However, from the point of view of the redundant data control 99, the cache memories function as a single common memory. The bandwidth between the redundant data control 99 and the cache transfer control 18 is the product of one disk's bandwidth and the number of drives in the array minus one. The disk controllers also include DMA controllers and independently perform direct

memory access transfers between the cache and each individual disk drive. In typical operation, the disks in the array are given instructions to transfer large blocks of data between the disks and the cache. The transfers are accomplished by the DMA controllers transferring smaller blocks. Once a DMA controller has completed an operation, it informs the control system, whereupon a new set of instruction is given to the DMA controller. A disk operation complete acknowledgment is generated by a disk array monitor after n-1 drives in the array of n drives have completed their operations and is sent to the Validator (described hereinbelow) for each DMA operation. Internal buffering within each drive buffers data while DMA is reprogrammed, thereby eliminating rotational latency delays. Contiguous data on the disk is thereby transferred to and from non-contiguous locations within the cache (see FIG. 2).

Cache to I/O transfer control block 18 is responsible for transfer of data between the cache and input/output data channels 19. Input/output data channels 19 convert incoming video/audio data into an internal format during record operations and convert back to the external video/audio format during playback operations. Any suitable video format may be employed.

In the illustrated embodiment, the array of disk drives 12 comprises 13 hard disk drives, each disk drive being a 1.6 GByte SCSI drive and the cache comprises 4 MBytes of memory per disk drive (for a total of 52 MBytes).

In the system according to the present invention, the algorithm for managing the cache and the organization of disk data are tightly coupled.

The illustrated embodiment of the invention stores both video and audio data on the same disk array.

Referring now to FIG. 2, a diagram of the arrangement of data according to the present invention, the data is organized on a disk array as a series of elements 20 one after another, each element comprising a video portion 22 containing active (picture) data for some fixed number of fields and an audio portion 24 containing a corresponding amount of audio data. Inactive (blanking) data is not stored in the disk array. Having video and audio separated in chunks is useful in that each may then be manipulated separately. For example, video could be played back while recording the audio, or the video could be played while skipping the audio altogether. The quantity of data in one element is determined by how often the disk array can perform seek operations while still enabling the average bandwidth to be sufficient to support the I/O data rate. The number of fields per element is typically constant, but is dependent on the video format. For example, in the illustrated embodiment of the invention, for D1 525 format signals, an element contains active video and audio data for twelve fields, whereas for D1 625 format signals, an element contains ten fields of active video and audio.

The cache 14 is of lesser length than the disk array and is organized as separate video and audio portions V_1 through V_m and A_1 through A_m respectively. Signal data is stored in the cache which is logically divided into a section of video blocks 25 and a section of audio blocks 27 wherein each video block may hold video data for a single disk element and each audio block may hold audio

data for a single disk element. Therefore, an element from the disk array corresponds to two blocks in the cache. In some cases a block may contain less than a complete element. The system
5 allows random access to data which is stored contiguously on the disk array. A portion of a disk array element may be transferred based on control commands received by the system from the user. In the illustrated embodiment, the cache is
10 of sufficient size to contain eight video blocks and eight audio blocks. The cache is ordered such that the material within the cache is in sequential order. The cache is organized as a circular structure. As the current position
15 crosses a cache block boundary, the block containing the oldest material is used to fetch future material for playback or reserved for input during record. This organization retains the most current video and audio output in the
20 cache to minimize disk requests during motion state changes.

The control system maintains a cache tag table for tracking information about data within individual cache data blocks. The cache tag
25 includes a tag field state for each video and audio field in the tag which indicates what data within the cache will be used for and with what step of the cache management process the data is currently engaged. The tag field state may be one of the
30 following states:

EMPTY - a field is empty and available for use
PLAY RESERVED - this field is reserved for
playback data but is empty
PLAY FILLING - data is being transferred from
35 the disks into this field

PLAY FILLED - all requested data is in the disk cache for this field and is ready for transfer to input/output data channels 19

5 RECORD FILLING - record data is transferring into this field from input/output data channels 19

RECORD FILLED - data in this field is ready to be written to disk

10 RECORD EMPTYING - the disk controllers are in the process of writing data from this field to the disks.

The various states available may be better understood in conjunction with FIG. 3 which illustrates the states and the transitions therebetween. From the EMPTY state, a field within the cache block
15 may change either to PLAY RESERVED state 62 or to RECORD FILLING state 64. Once the entry is in the PLAY RESERVED state, as data is requested from the disk array, a transition between PLAY RESERVED and PLAY FILLING 66 will occur. The PLAY FILLING state
20 transitions to the PLAY FILLED state 68 once all requested data is in the disk cache for this block and is ready for transfer to I/O data channels 19. From the PLAY FILLED state, the tag's field state may then transition to PLAY RESERVED, to RECORD
25 FILLING or back to the EMPTY state as the block is needed for future playback or record operations.

The RECORD FILLED state 70 is reached from RECORD FILLING state 64 when all requested data is in the disk cache and ready to be written to disk.
30 From the RECORD FILLED state, transition may be made to the RECORD EMPTYING state 72 when the disk processors are in the process of writing data to the disk. Once the data is written to the disk, the state is changed to PLAY FILLED, allowing playback
35 or freeing the cache block for further use.

The cache tag also includes the playback starting and ending fields, the field corresponding to the first field of the disk element mapped to the tag, the cache address of each field of the tag, the disk array containing the data, and the late disk, identifying which was the last disk to respond.

It is the responsibility of the control system 16 (FIG. 1) to retrieve data from the disk array before the data is needed for playback or to provide temporary storage of data during a record operation until such time as the data can be written to the disk array. In each case, the control system must predict what will be needed in the future, this prediction being a function of the current position within the data, the system motion state and the play speed. Current position is the field which is either currently being provided as input or which is being output. The general goal of cache management is to have enough material in the cache surrounding the current position to satisfy any action that the operator may make in changing the motion state of the system.

The control system is responsible for executing high-level commands, providing disk read and write commands to the disk array and monitoring each drive in the array to determine completion of the data transfer.

In playback mode, the disk array must read information in advance of when it will actually be output and store the information in the disk cache. The information in the disk cache may then be transferred to video and audio outputs in real time. In record mode, incoming video and audio data are transferred to the disk cache in real time, and the

disk array is instructed to write the cached information an element at a time.

Cache management is divided into three subsystems that deal with operation of the disk cache and the disk array: Predictor; Validator; and Query. Predictor determines what material should be in the cache based on the current position within the data and the current motion state. The motion state may be one of Stopped, Playing, Recording, or Shuttling, which are described hereinbelow with reference to Table I. Predictor also translates between playlist space and reel space. Reel space is a sequential series of fields stored on a disk array. A reel may span more than one disk array or several reels may share a disk array. The system contains random-access playback capabilities which are provided to the user through the use of playlists. The system can access material scattered randomly across the disk array in addition to the standard linear manner of playback. Playlist space comprises several segments, a segment being a set of one or more contiguous output fields. Each playlist may identify different portions of a reel. A field that is defined in terms of playlist space is designated as an offset of a number of fields relative from the first field in the playlist. Predictor will generate requests for material, e.g., video data or audio data, based on the current motion state, the contiguous desired material that is currently within the cache and space available in the cache. Transfers may include playback (transfers from the disk array to the cache) and record (transfers from the cache into the disk array). Predictor is performed once each field.

Validator updates the state of the cache according to the acknowledgments received from the disk array. These acknowledgments indicate what data has been transferred to and from the disk array. Validator is also responsible for indicating that sufficient material has been transferred into or out of the cache to allow further commands to be processed.

Query facilitates transfer of data between the cache and input/output data channels 19. It does this by providing cache addresses for location of video and audio materials that are needed for playback or providing the addresses necessary to indicate where newly recorded material will be placed within the cache prior to transfer to the disk array.

The function of the Predictor is presented as a state table in Table I. The state table consists of the previous state along the horizontal axis and a new state along the vertical axis. Each box of Table I represents the procedure followed at the state transition. At any one time, the previous state or new state may comprise one of the following states: Initialize; Stopped; Playing; Recording; or Shuttling. Initialize is the state which occurs when the recorder is first started. The Stopped state takes on two modes: steady state and go to request. In the steady state mode, Stopped is the state wherein the current position is not changing. In the case of a go to, Stopped state processes a request from the user to output a field other than the current position. The system processes this go to request by assessing the playlist field requested relative to the current contents of the cache and only the data needed to center the

request field within the cache is requested from the disk array. In the Playing state, the recorder retrieves data from the disk arrays and supplies the retrieved data as output through input/output data channels 19. The Recording state involves transferring received data from input/output data channels 19 to the disk arrays. Shuttling is a state wherein the recorder plays data.

10

Previous State

| New State | Initialize | Stopped | Playing | Recording | Shuttling |
|-----------|------------|-------------------------------|--|--|---|
| | Initialize | Invalid transition | Invalid transition | Invalid transition | Invalid transition |
| | Stopped | Set cache busy Build Cache | Crossed Elem Time to Fetch Build Cache | Set cache busy Write Residue Build Cache | Set cache busy Build Cache |
| | Playing | Invalid transition | Play | Set cache busy Write Residue Calculate Gap Play | Set cache busy Calculate Gap Play |
| | Recording | Invalid transition | Record | Set Tracks Record | Record |
| | Shuttling | Invalid transition | Shuttle | Invalid transition | Shuttle |

TABLE 1. Predictor State Transition Table

With reference now to Table I, the various operations of the Predictor will be described. A number of processes are set out in Table I and are defined as follows:

Set cache busy - sets the state of the cache status to indicate to the system that the cache management process is not available for command input.

Build Cache - performs an analysis of the cache to determine the usable contiguous material within the cache, the amount of cache space that is available, the size and distribution within disk elements of desired segment and playlist fields. Build Cache also determines where new material will be placed within the cache, prioritizes and generates disk requests.

Crossed Elem(ent) - determines if the current position within the play or record data has crossed a cache element boundary. If a boundary has not been crossed, no further steps are taken.

Time to Fetch - determines if the amount of time that would be necessary to fetch the necessary material exceeds the limits of the system and if so, sets the state of the cache to busy.

Play - maintains a centered cache as the current position moves through the cache. The cache is centered when the current position has equal amounts of forward and reverse play data in the cache on either side of the current position. Play optimizes and generates disk read requests to transfer data from the disk array that may be needed to fill the cache.

Record - maintains a centered cache as input source data is recorded into the cache. Optimizes and generates disk requests to write recorded material from the cache to the disk array.

Shuttle - determines the expected field based on shuttle speed and change of direction. Generates disk read requests to fill the cache with data from the disk array.

Write Residue - determines if the system has completed the recording of source material into the cache. If the recording is complete, a disk write

request is issued to transfer material from the cache to the disk array.

5 Calculate Gap - determines the portion of the cache which will not be usable because of transfer time constraints and other system constraints.

Set Tracks - determines whether, for example, video tracks, audio tracks or both are being recorded.

10 Invalid transition denotes that the state is not properly attainable during normal operation of the recorder.

The rules which the Predictor applies in operation are:

15 In forward play, the Predictor reserves cache blocks as far forward in time as needed to center the current output field in the cache. Future material will occupy up to five cache blocks ahead of the block containing the current output; likewise, previously played material will occupy about three
20 blocks prior to the block containing the current output. Blocks previously played are maintained until needed for future playback material. This assures that material needed in the cache for state transitions (i.e., Play to Stopped) is available
25 without necessity of a disk request. For reverse play, this same algorithm is followed; however, a backward biased cache is built, again maintaining a centered output with approximately five blocks proceeding the block containing the current output.

30 The system is in the Record state when any or all tracks are recording. The procedure followed during Record is one of allocation of cache blocks for use by the input data. As a single block is filled by the input data, the cache manager issues
35 disk requests for the disk array to transfer data

from the cache to the disk array. The cache manager then frees the next block as available for input. In order to minimize disk requests during state transitions, the material which has been previously recorded and transferred to disk is maintained in the cache. This assures that if a transition from Record to Play is made before the recorded material has filled the entire cache, the material contained in the cache prior to the record is still available. Likewise, during a transition from Record to Play or Stopped, the material which was just recorded remains in the cache requiring that only the material ahead of the recorded data needs to be fetched from the disk array.

In the Shuttle state, single fields are placed into the cache blocks. The field which will be requested from the disk array is determined by the current output position plus a shuttle constant times the shuttle speed, e.g., with a current output field of 100 and a speed of 4 (unity is 1) and a constant of 3, the next field to be placed in the cache block adjacent to the block containing the current output is field 112. Again, the previously output data is retained in the cache to allow for change in shuttle direction, e.g., reverse to forward.

In the Stopped state, the cache is built around the current output position. The current output position is the output field requested by the user. Depending on the previous state of the cache (particularly the current output position) more or fewer disk reads will have to be generated to center the current output position. If the current output position is not contained within the cache, the entire cache must be rebuilt from the

disk array. This current output position must be fetched from the disk array immediately. Therefore, the material requested from the disk array follows the pattern: request current output position,
5 request data needed ahead of the current output position, request data needed behind the current output position. If the current output position is contained in the cache, the disk requests needed to center this position are generated. An assessment
10 of the cache is made to assure that the minimum number of fields are requested from the disk array.

The Validator process is responsible for maintaining the cache state in response to acknowledgment of data transfers to and from the
15 disk arrays. Referring now to FIG. 4, a flow chart of the Validator process steps performed in cache management, the Validator process begins with Receive acknowledgements (step 30), wherein acknowledgement signals are received from the disk
20 array indicating whether an amount of material has been transferred to the cache from the disk array when in play mode or information has been transferred from the cache to the disk array when in record mode. Descramble (step 36) determines
25 what type of material was received from the disk, video or audio or both, and translates the received information into cache fields. Next, the cache tag field states are updated (step 38) which sets the state of the cache to indicate which blocks of
30 video and audio are now available for playback or record. The various states which may be assumed by cache blocks are described hereinabove with reference to FIG. 3. Following step 38, decision block 40 determines whether an acknowledge was
35 received at step 30 and whether the system is

currently in a play state. If the decision is YES, then step 42, determine last disk responded, is performed. The determine last disk responded step updates the state of the cache to indicate which disk within the array was the last to respond relative to the most recently received acknowledgment. After step 42, the Validator process continues directly following decision block 40. If the result of decision block 40 is NO, then step 42 is not performed. Next, decision block 44 determines whether material is available. Material available indicates whether material has been transferred into or out of the cache sufficient to allow further cache data transfer commands to be processed. If material is available, then step 46 updates the system cache state to an available status, which indicates that further cache commands may be processed. After step 46 or following decision block 44 if a NO results from the decision, the Validator process is completed.

FIG. 5 is a flow chart of the Query process. Query is called by the system control to either determine where to place newly received recorded material in the cache, or to determine the location within the cache of material needed for playback. Query begins with decision block 48, which determines whether a record state field is being requested. If a record field is requested, then a search of the cache is made (block 50) in reel space to determine where the record data will be placed in the cache. Following step 50, the Query process is complete. However, if the result of decision block 48 is NO, then decision block 52 is performed to determine whether the cache request is for either the Playing or Stopped states. If the

result is YES, then a search is made in the cache (block 54) for appropriate material within the cache. Again, Query is complete after block 54.

5 If the result of decision block 52 is NO, then that system is in Shuttling state and step 56 determines a range of cache locations that would be acceptable to provide appropriate shuttle playback data. Following step 56, the cache is searched in step 58 to determine the position
10 within the cache of data that meets the range criteria determined in step 56. The Query operation is then completed.

As mentioned hereinabove, the system provides redundancy in the disk array which enables
15 playback to occur even though one drive may not have responded. Such an operation would not be appropriate in record mode since the redundant data must be transferred reliably to be of any use in playback and, hence, the whole set of disks must
20 respond before a record operation is completed. Employing such redundancy enables the system to provide playback data without waiting for the slowest disk of an array to respond with data or without waiting to determine that a disk drive has
25 failed. Therefore, the response of the system is not dictated by the slowest drive, providing a speed improvement on playback. Further, it is not necessary to determine that a drive has actually failed before employing redundancy to reconstruct
30 playback data; the reconstruction is automatically performed after n-1 drives respond with data. Since the identity of the non-responding (or last to respond) drive is known, reconstruction of data is accomplished by exclusive ORing the data from
35 the n-1 drives thereby producing the missing data.

The resultant data is then inserted at the position where the data from the non-responding or late drive would have been. Of course, if the non-responding drive is the one that contains the redundant data bit, no reconstruction is necessary. Thus, the recorder provides redundant storage of video and audio data.

Therefore, a system and method for managing a disk cache for a digital video recorder have been shown and described wherein a disk array is managed to enable continuous data to be written to or supplied from a series of non-continuous access disks wherein the disks are more likely to be operating in a burst mode rather than in a continuous data stream mode wherein the burst rate of the disk array (the rate at which individual groups of data are transferred to the array) is higher than the continuous data stream rate. The term "continuous data" may be understood to include data transferred external to or received external from the recorder at a uniform clock rate, while non-continuous access storage may include the transfer of data at non-uniform clock rates, sometimes at a higher rate than the "continuous" uniform rate, but other times at a much lower (e.g. zero) rate.

While the particular embodiment employs video and audio data, other information may also be used, for example, video and key data or video and key/depth data, wherein depth data may define layering priority of corresponding video data. The video portion of an element may also contain both a video signal and corresponding key data. Further, multiple disk reorders may be coupled together to provide multiple disk arrays to allow more record

and playback storage time or to allow more channels
of data to be recorded concurrently under control
of one control system. Therefore, it will be
appreciated that the invention is not restricted to
5 the particular embodiment that has been described,
and that variations may be made therein without
departing from the scope of the invention as defined
in the appended claims and equivalents thereof.

10

15

20

25

30

35

Claims

1. Apparatus for recording and playing continuous data comprising:

an array (12) of non-continuous access storage devices for storing data;

a cache (14) for either receiving continuous data from an external device and storing said received continuous data until said continuous data can be transferred in discrete portions to said array of storage devices or retrieving discrete portions of data from said array of storage devices and storing said retrieved data until said retrieved data is transferred as continuous data to an external device; and

computer means (16, 18) for directing operation of said cache to provide interchange of data between said array of storage devices and an external device.

2. Apparatus according to claim 1, wherein said cache comprises:

memory for holding said continuous data from an external device and for holding retrieved data from said array of non-continuous access storage devices.

3. Apparatus according to claim 1, for recording and playing video data, wherein said cache comprises a plurality of video data storage blocks (25) for storing video data.

4. Apparatus according to claim 1, for recording and playing video data and associated audio data, wherein said cache comprises a plurality of video data storage blocks (25) for storing video

data and a plurality of audio data storage blocks
(27) for storing audio data.

5 5. Apparatus according to claim 1, wherein
said computer means further comprises program means
for managing organized exchange of data between an
external device and said cache and between said
cache and said array of non-continuous access
storage devices.

10 6. Apparatus according to claim 1, wherein
said array of non-continuous access storage devices
comprises an array of disk drives.

15 7. Apparatus according to claim 1,
further comprising redundant data storage means
(17, 99) for recording redundant data and enabling
playback of valid data even when less than the
entire array of non-continuous access storage
20 devices has provided playback data.

 8. A method for managing a cache of
continuous information to transfer the continuous
information between an external device and an array
25 of non-continuous access storage devices, the
method comprising the steps of:

 predicting what information should be in
the cache;

 reserving space in the cache based on said
30 predicting step;

 establishing a cache tag table
representing the status of information requested
for input to and output from the cache; and

 transferring information between said
35 reserved cache space and an external device and

between said reserved cache space and said array of non-continuous access storage devices.

5 9. The method according to claim 8, wherein said step of predicting what information should be in the cache comprises the substeps of:

10 if in a playback mode, indicating that cache space should be reserved for information which will be needed from the array of non-continuous access data storage devices, said indicating including marking portions of said cache space as reserved for play, and

15 if in a record mode, indicating that cache space should be reserved for information which will be arriving from an external device, said indicating including marking portions of said cache space as reserved for record.

20 10. The method according to claim 8, further comprising the steps of:

optimizing cache space by retaining previously used information in the cache; and updating the state of the cache tag table as necessary,

25 whereby smooth state transitions may be facilitated.

30 11. The method according to claim 10, wherein said step of optimizing cache space by retaining previously used information comprises the substeps of:

35 maintaining a circular cache whereby cache space containing information least likely to be accessed is made available for storage of information predicted to be needed in the future;

in play mode, centering the current output information in the cache by use of the circular cache;

5 in play mode, marking the oldest portion of the cache available for transfer from the array of non-continuous access storage devices only as current output moves forward or backward away from cache center; and

10 in record mode, marking cache portions which contain recorded data available for transfer to said array of non-continuous access storage devices, initiating said transfer, marking the recorded data containing cache portions available for playback and marking as available for record
15 the portions of the cache containing information least likely to be recorded or played back from the cache.

20 12. The method according to claim 8, wherein said step of transferring information between said reserved cache space and an external device and between said reserved cache space and said array of non-continuous access storage devices comprises:

25 if in a play mode, reading play information from said array of non-continuous access storage devices into portions of the cache marked as available to be filled with play information, and

30 if in a record mode, writing record information to said array of non-continuous access storage devices from portions of the cache which are marked as containing newly received unwritten information.

13. The method according to claim 12, wherein said step of reading play information comprises the substeps of:

5 selecting a reserved for play portion in the cache;

 requesting an operation to read information into the selected cache portion from said array of non-continuous access storage devices; and

10 changing the cache tag table state of the selected portion to indicate a disk request has been generated, and when the read operation is completed, changing the cache tag table state of said selected portion to indicate information
15 is available for output.

14. The method according to claim 12, wherein said step of writing record information comprises the substeps of:

20 detecting a cache portion which has been fully recorded;

 requesting a write operation to write information from the portion to said array of non-continuous access storage devices;

25 changing the cache tag table state of the portion to indicate a write operation was requested; and

30 changing the cache tag table state of the selected portion to an available for output state when the write operation is complete.

15. The method according to claim 8, wherein said continuous information comprises video.

16. The method according to claim 8, wherein said continuous information comprises video and audio.

5 17. A video disk recorder comprising;
 an array (12) of n non-continuous access
 storage devices for storing video and audio data,
 one of said storage devices storing redundant data;
 a memory cache (14) for either receiving
10 continuous data from an external device and storing
 the received continuous data until the continuous
 data can be transferred in discrete portions to the
 array of storage devices or retrieving discrete
 portions of data from the array of storage devices
15 and storing the retrieved data until the retrieved
 data is transferred as continuous data to an
 external device; and
 computer means (16, 18) for directing
 operation of said memory cache to provide
20 interchange of data between the array of storage
 devices and an external device.

 18. Apparatus according to claim 17,
 wherein said computer means comprises means (99)
25 for retrieving video and audio data from the array
 of storage devices after n-1 of the storage devices
 have responded to a request for data retrieval.

 19. A digital video recorder
30 according to claim 17, wherein said memory cache
 is organized as a plurality of blocks (25) and a
 plurality of said memory cache blocks are adapted
 to store video data.

35

20. A digital video recorder according to claim 17, wherein said memory cache is organized as a plurality of blocks (25) and a plurality of said memory cache blocks are adapted to store multiple video data.

21. A digital video recorder according to claim 17, for recording and playing both video data and audio data, wherein said memory cache is organized as a plurality of blocks and a first plurality of said memory cache blocks (25) are adapted to store video data and a second plurality of said memory cache blocks (27) are adapted to store audio data.

22. A method for recording and playing continuous data comprising:

- receiving continuous data from an external device;
- storing the received continuous data in a cache (14);
- transferring the continuous data in discrete portions to an array (12) of non-continuous access storage devices thereby recording the data;
- receiving non-continuous data from the array of non-continuous access storage devices;
- storing the received non-continuous data in said cache; and
- transferring the non-continuous data to an external device in a continuous manner thereby playing the data.

23. The method according to claim 22, wherein the received continuous data comprises video data.

24. The method according to claim 22, wherein the received continuous data comprises video data and audio data.

5 25. A method for retrieving information from an array of n data storage devices, said array including one redundant storage device, comprising the steps of:

10 requesting information from the array of data storage devices;

 retrieving data from the first $n-1$ data storage devices to respond to the request for information; and

15 deriving the information based on the retrieved data, thereby enabling information retrieval without waiting for the response of the last data storage device.

20 26. The method according to claim 25, wherein the information is stored on the array of storage devices by a storage method comprising the steps of:

 storing information on $n-1$ storage devices;

25 determining redundancy data from the information; and

 storing the determined redundancy data on the n th data storage device;

30 27. The method according to claim 26, wherein said step of determining redundancy data comprises exclusive ORing portions of information stored on each of the $n-1$ storage devices.

35

28. Apparatus for playing and recording continuous data substantially as herein described with reference to and as shown in the accompanying drawings.

5 29. A method substantially as herein described with reference to and as shown in the accompanying drawings.

- 30 -

Relevant Technical Fields

Search Examiner
MR P R SLATER

(i) UK Cl (Ed.L) G5R(GB), G4A(AKB1)

(ii) Int Cl (Ed.5) G11B, G06F

Date of completion of Search
19 November 1993

Databases (see below)

(i) UK Patent Office collections of GB, EP, WO and US patent specifications.

Documents considered relevant following a search in respect of Claims :-
1-24

(ii) ONLINE DATABASES:WPI

Categories of documents

- | | |
|--|---|
| <p>X: Document indicating lack of novelty or of inventive step.</p> <p>Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.</p> <p>A: Document indicating technological background and/or state of the art.</p> | <p>P: Document published on or after the declared priority date but before the filing date of the present application.</p> <p>E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.</p> <p>&: Member of the same patent family; corresponding document.</p> |
|--|---|

| Category | Identity of document and relevant passages | | Relevant to claim(s) |
|----------|--|----------------|----------------------|
| X | EP 0508441 A2 | (MITSUBISHI) | 1,17,22 at least |
| X | EP 0507552 A2 | (IBM) | 1,17,22 at least |
| X | EP 0503768 A1 | (IBM) | 1,17,22 at least |
| X | EP 0493984 A2 | (ARRAY TECH) | 1,17,22 at least |
| X | WO 92/01988 A1 | (CAB-TEK) | 1,17,22 at least |
| X | WO 90/00280 A1 | (STORAGE TECH) | 1,17,22 at least |

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).